PRECISION MANIPULATION WITH COOPERATIVE ROBOTS

Ashley Stroupe, Terry Huntsberger, Avi Okon, Hrand Aghazarian

Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive, Pasadena, CA, 91109
{Ashley.Stroupe, Terry.Huntsberger, Avi.Okon, Hrand.Aghazarian}@jpl.nasa.gov

Abstract: We present a cooperative approach to robotic precision manipulation tasks in the context of autonomous robotic construction. Precision manipulation requires a firm grasp, which constraints the team to rigidly maintain formation during transport and manipulation. A leader/follower approach with force sensing to provide relative formation information and vision to provide team position relative to construction components is applied. Our approach demonstrates successful, reliable performance of a construction task requiring cooperative transport and placement of structure components. Qualitative and quantitative performance results are provided.

Key words: Cooperative Robots, Cooperative Transport, Robot Construction

1. Introduction

The 2004 NASA vision for space exploration calls for a sustainable presence in space, beginning with a human return to the Moon in 2020 (NASA, 2004). A sustainable robotic or human presence requires deploying and maintaining infrastructure to provide power, living quarters, and resource acquisition and utilization; deployment must be autonomous for safety and reliability. Space operation places constraints on rover power, computing, communication, and mass. JPL's Planetary Robotics Lab (PRL) is developing autonomous technologies to perform construction related tasks under these constraints. One focus area is cooperative transport and precision manipulation of large rigid components over natural terrain using fused sensor information from both robots. We present details and quantitative results of our approach, an extension of previous work (Stroupe, et al., 2004) and done under the CAMPOUT architecture (Huntsberger, et al., 2003).

This work addresses several challenges of cooperative transport and precision manipulation. Precision manipulation requires a rigid grasp, which places a hard constraint on the relative rover formation that must be accommodated, even though the rovers cannot directly observe their relative poses. Additionally, rovers must jointly select appropriate actions based on all available sensor information. Lastly, rovers cannot act on independent sensor information, but must fuse information to move jointly; the methods for fusing information must be determined.

2. Background and Related Work

Precision manipulation of large components requires rovers to have a rigid grasp and therefore to remain in a fixed relative formation, even if rovers cannot directly obtain each other's relative position. Cooperative transport has primarily focused on pushing on flat floors, typically relying on direct observation or communication of relative position (Parker 1994, Qingguo and Payandeh, 2003; Rus, et al., 1995). In some cases, robots use forces imparted by the object to communicate, but robots adjust position to obtain desired forces (Brown and Jennings, 1995). These approaches are not applicable to precision manipulation in three-dimensions, rigid formations, or operation in natural terrain. Most precision cooperative manipulation using force feedback for implicit communication is for fixed-base manipulators (Mukaiyama, et al., 1996), and is not applicable for construction. Formation following typically uses potential fields, which require observation or communication of relative pose and do not accommodate hard constraints (Balch and Arkin, 1998; Carpin and Parker, 2002; Desay, et al., 1999). Most similar to our task is Omnimate (Borenstein, 2000) which uses a compliant linkage between two mobile platforms; Omnimate compensates for uneven floors and moderate wheel-slippage by controlling wheel velocity based on observed angular difference between the expected and observed lines of contact between the platforms. To date, cooperative transport in rigid formation is limited to JPL's Robot Work Crew (Huntsberger, et al. 2004, Trebi-Ollenu, et al., 2002), and mobile precision manipulation has not yet been demonstrated.

3. Cooperative Transport / Manipulation

3.1 Task Domain and Description

Two rovers, heterogeneous in size and arm configuration, are equipped with a forward-facing stereo camera pair and a three-axis force-torque sensor on the arm's wrist. The team can cooperatively carry long beams that are stacked and interlocked if positioned accurately. Each beam has a grasping point at either end; the gripper lower finger passes through the grasp point and the upper fingers close on the top of the beam. Rovers obtain beam location from the stereo vision of three fiducials on each end. The test environment (Figure 1) is a PRL sand pit that simulates natural terrain.

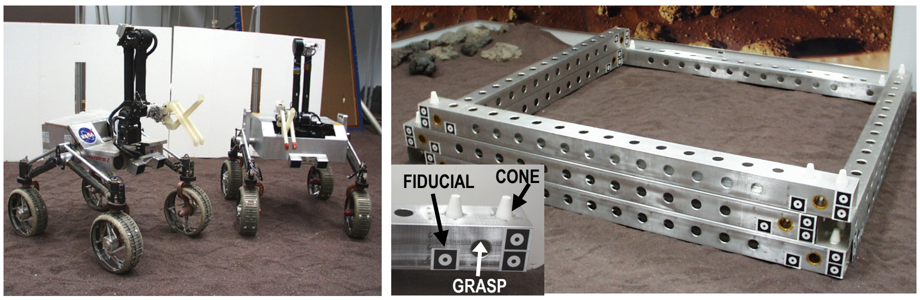

Figure 1. Left: RCC Rovers SRR (left) and SRR2K (right) are heterogeneous; note SRR2K's gripper is not in line the rover body center. *Right*: Structure composed of stacked beams; note each beam is alternately offset laterally from the one it rests upon. The inset shows the grasping point, the fiducials and the interlocking cones on one beam.

Tasks are in the context of building a four-sided structure by stacking beams. The team must cooperatively transport beams from storage to the structure and then align at the structure while remaining in formation to prevent dangerous forces on rover arms. Once aligned, the team must precisely place the beam in the structure while keeping arms at appropriate relative positions. Robots return to storage and repeat the process to build the structure. Current work demonstrates end-to-end acquisition and placement of a single beam; results here focus on the cooperative aspects.

3.2 Cooperative Behaviors

3.2.1 Aligning with the Structure

Rovers position themselves relative to the structure based on stereo vision of the fiducials on the beams already in the structure. To place a beam into the structure, the rover arms must maintain separation; the rigid grasp cannot accommodate arbitrary offsets in rover position. Thus, team alignment with the structure must be precise (within 1 cm). Slippage, errors in drive kinematics, and errors in estimated structure location lead to alignment errors. Vision errors are minimized by validating observations against the model of fiducials on the beam and repeating observations as necessary. Additionally, the vision is calibrated relative to the arm to offset any kinematics errors in the arm control. The team reports any unrecoverable errors to allow corrective intervention.

To reduce the effects of errors, the algorithm uses an iterative process of adjusting range, lateral offset, and heading, as illustrated in the state machine of Figure 2 and the diagrams of Figure 3. To ensure rovers are executing parallel motions, they share data and the leader selects an action and sends to the follower; roles can be dynamically changed. As position is refined, visual information becomes more accurate and errors from slip in previous motions are corrected. If an iteration completes without corrections, the team is properly aligned. To maximize visual accuracy, the team iteratively aligns directly in front of structure fiducials first, and then iteratively aligns the lateral offset appropriate for placing the beam. The leader selects an action (movement magnitude and direction) using fused information from both robots, if both see the structure, or the individual information from either robot, if only one robot sees the structure.

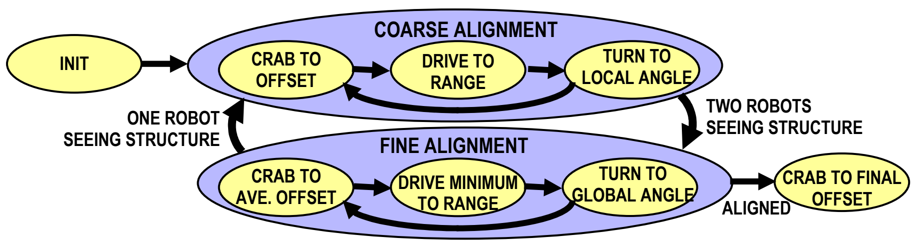

Figure 2. Simplified finite state machine representing the alignment process, with different modes for action selection depending on the amount of data available.

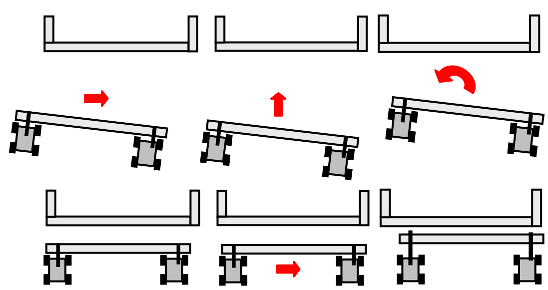

Figure 3. Top: Iterative alignment process of crab (left), drive (center), and Ackermann turn (right). *Bottom*: Aligned team (left) crabs to offset beam for placement (center) and extend arms to place beam (right).

When only one robot sees the structure, the team must move the team so as to allow the other rover to also see the structure. This involves attempting to position the other rover at an appropriate range and lateral position. Due to the linear formation constraint imposed by the beam, if one rover is approximately aligned, the other rover should be near alignment and able to see the structure. Thus, the observing rover approximately aligns itself to bring the other rover into position. Driving (forward or backward) and crabbing (lateral driving) will bring the observing rover close to desired range (X_D) and zero arm lateral offset by moving the team those distances;

this is shown in (1), where (X, Y) is the observed structure position and M_Y is the manipulator arm lateral offset. To bring the team parallel, the rovers execute an Ackermann based on the observing rover's estimate of its heading relative to the structure; this is determined using the difference in range to the fiducials on the left and the fiducials on the right (ΔX_T) and the known lateral distance between fiducials (ΔY_T). The rover turns to counter this angle as in (2).

$$D = \begin{cases} X - X_D \\ Y - M_Y \end{cases} \tag{1}$$

$$A = -\tan^{-1}\left(\Delta X_T, \Delta Y_T\right) \tag{2}$$

When both robots see the structure (during fine alignment and during the final lateral offset), information is fused differently for different conditions and different motions. Driving brings one member of the team to the correct range. The team selects to execute the shortest drive (X) computed by (1) to prevent overshooting desired range. Crabbing aligns the rovers laterally with the structure. The team executes a crab in magnitude to satisfy both rovers: the average offset. If the rovers are in a linear formation and aligned with the structure, the average Y offset from (1) will bring both rovers into position. To bring the rovers to the same range (aligning the team and beam parallel to the structure), an Ackermann turn is done. The turn is computed as in (3), using the difference in range (ΔX) for the two robots and the distance between beam grasping points (ΔY_G); the further robot drives forward while the closer robot drives backward, equalizing the distance. There is one exception: if in the previous step only one rover saw the structure and the team anticipated turning based on local rather than global information, the follower's range (required for determining global angle) is not available, but instead the follower's local angle was shared. In this case, the turn is the average of both local angles from (2) to approximately satisfy both. Rather than attempting to robustly identify this situation on both rovers and resend the proper data, this approach saves additional communication and ensures the rovers stay in parallel states. Turns too small to execute (less than one degree) are approximated by a drive, with each rover driving independently to the correct range from (1); these small offsets can be tolerated in forward driving.

$$A = \tan^{-1}\left(\Delta X, \Delta Y_G\right) \tag{3}$$

If both robots lose sight of the structure, they move forward in increments of 5 cm until at least one regains observation or until a timeout is reached.

3.2.2 Driving in Rigid Formation

As described in 3.2.1, the team performs three motions in formation: drive, crab, and Ackermann turn. To keep the beam straight and prevent large stresses on the arms, the team must remain at the desired lateral separation and in-line with each other. Keeping formation allows the team to achieve the correct relative position to place the component in the structure. The arms can only accommodate small lateral offsets during component placement without inducing large forces, thus accurate rover positioning is essential. By implementing motions determined by the leader (rather than in a distributed manner), cooperative moves are always identical. Synchronizing cooperative moves reduces (but does not eliminate) errors due to time offsets. Leader-controlled motion and synchronization do not mitigate initial formation errors.

Empirical data has indicated correlations between torque about the vertical axis and rover alignment and between force along the horizontal axis and rover separation; the team detects errors in relative formation using forces and torques imparted through the beam. Once a formation error is detected it can be corrected as shown in Figure 4. By allowing the follower to adjust its velocity, it can eliminate formation errors. For driving, the follower speeds up if torque indicates it is behind and slows down if it is ahead, Figure 4 illustrates the case where the follower is on the right. The relationship between torque and force and velocity correction is dependent on direction of motion. Currently, only velocity control is used to adjust formation; steering adjustments to account for off-axis formation errors is in development. To ensure forces and torques do not build up before velocity control can compensate, the follower starts first (the leader delays 2 seconds after synchronization). To prevent overreaction to sensor noise, the sensor is sampled at 20 times the action selection rate and averaged.

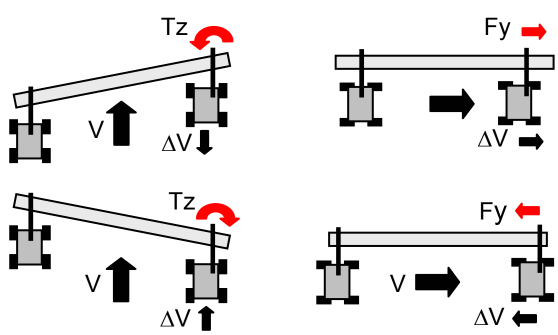

Figure 4. Left: Follower adjusts velocity to eliminate torque. The follower slows down to reduce counterclockwise torque and recover alignment (top) or speeds up to reduce clockwise torque and recover linear alignment (bottom). *Right*: Follower adjusts crabbing velocity to eliminate force. Follower speeds up to increase spacing and reduce positive force (top) or slows to decrease spacing and reduce negative force (bottom).

The velocity controllers are PI controllers designed to maximize response time, minimize oscillation (particularly at the end of moves), and stop both rovers upon failure. Parameters for control were empirically determined based on observed performance. (4) shows the general controller, where V_B is the base velocity, E is the error term, ΣE is total accumulated error, and ΣE_{MAX} is the magnitude limit on the accumulated error. (5) shows drive velocity error, where T_Z is torque about the vertical axis T_{Z0} is the reference torque about the vertical axis, T_{FY} is torque about the vertical axis due to lateral forces rather than angular misalignment of the beam. Drive control parameters are $K_P = 0.4$, $K_I = 0.0$, and $\Sigma E_{MAX} = 0.0$. (6) shows the error for crabbing velocity, where F_Y is lateral force and F_{Y0} is reference lateral force. Control parameters for crabbing are $K_P = 0.6$, $K_I = 0.05$, and $\Sigma E_{MAX} = 6.0$. The single-step change in velocity is limited by ΔV_{MAX} as in (7).

$$V_{NEW} = V_B + K_P E + K_I \Sigma E \tag{4}$$

$$E_D = T_Z - T_{Z0} - T_{FY} \tag{5}$$

$$E_C = F_Y - F_{Y0} \tag{6}$$

$$V = \begin{cases} V_{LAST} + \Delta V_{MAX} & \text{if } V_{NEW} - V_{LAST} > \Delta V_{MAX} \\ V_{LAST} - \Delta V_{MAX} & \text{if } V_{NEW} - V_{LAST} < \Delta V_{MAX} \\ V_{NEW} & \text{otherwise} \end{cases} \tag{7}$$

While only the follower uses force-torque feedback to control velocity, both rovers monitor forces and torques. If either rover detects a sustained force or torque larger above a threshold, failure is detected and the rover stops driving. This quickly increases force and torque on the other (still moving) robot past threshold, so that it also detects a failure and stops. These thresholds are set based on empirical data. The maximum allowed torque magnitude about the vertical axis is 4.0 N-m and the maximum allowed lateral force magnitude is 44.6 N.

4. Experimental Results

4.1 Cooperative Transport with Adaptive Velocity

To demonstrate the team's ability to remain in formation and keep arm forces and torques within bounds, a series of tests were performed. For these experiments, the rovers start in formation with the follower on the right. Forces and torques minimized and the reference torques and forces are set to the initial values. A drive or crab command is sent and behavior is observed. Between repetitions of an experiment, the rovers are repositioned to reduce forces and torques, but the reference force/torque is not reset; this investigates ability to accommodate small initial offsets in formation. Force-torque profiles with velocity control and fixed velocity are compared.

Table I and Table II summarize performance. The resulting mean and standard deviation of the force or torque during nominal driving is reported (from after initial reaction time to initiation of deceleration).

Table I: Forward Drive Torques (N-m)

Distance	Velocity Control	Fixed Velocity, No Delay	Fixed Velocity, Delay
300 cm	-0.04 ± 0.16	-0.43 ± 0.25	-2.09 ± 1.02
	-0.05 ± 0.22	-0.26 ± 0.43	-2.22 ±0.80
	-0.04 ± 0.21	-0.95 ± 0.43	Failure
	0.02 ± 0.20	-0.06 ± 0.69	-2.03 ± 0.79
	-0.03 ± 0.20	-1.76 ± 0.79	-2.30 ± 0.81

Table II: Crab Drive Forces (N)

Distance	Velocity Control	Fixed Velocity, No Delay	Fixed Velocity, Delay
80 cm	-3.76 ± 7.16	3.91 ± 8.67	Failed
	0.05 ± 5.38	-1.48 ± 8.67	-14.04 ± 7.59
	-0.44 ± 4.59	-5.54 ± 6.66	-7.30 ± 9.42
-80 cm	1.61 ± 6.15	-6.73 ± 12.38	5.27 ± 9.90
	-0.02 ± 3.93	6.98 ± 11.06	Failure
	-0.78 ± 4.45	15.56 ± 9.86	-3.56 ±9.34

Results show velocity control keeps forces and torques low (within bounds) and with low variance despite initial offsets; fixed velocity shows higher torques (approximately 19 times higher mean) with a high bias and high variance. Fixed velocity may fail, particularly with start offsets, and does fail or nears failure in nearly half of the crabbing experiments. Torque profiles for drives with and without velocity control are compared in Figure 5. Torque remains near zero and has little variance while under adaptive velocity control despite initial errors, while torque has large bias and large variance without velocity control. By moving first, the follower can immediately respond to any change in forces and torques.

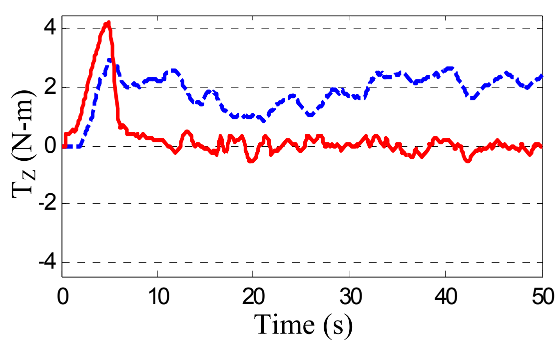

Figure 5. Torque profiles of a 300 cm drive with velocity control and a delay (solid) and without velocity control and no delay (dotted). Note high torques (near failure), large variance, and large bias without control.

Crabbing demonstrates similar results (Figure 6); velocity control keeps forces near zero with low variance but fixed velocity nears failure and has high variance.

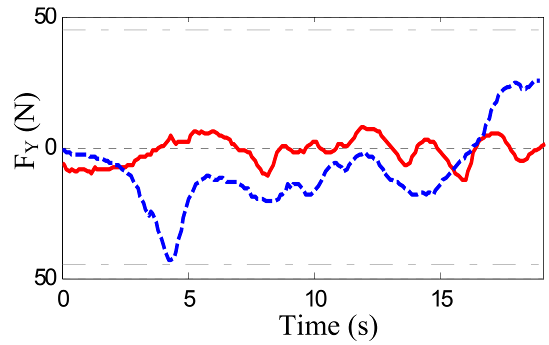

Figure 6. Force profiles of an -80 cm crab with the follower on the right. Note high forces and large variance without control (dotted) compared to velocity control (solid). Force bounds shown dashed.

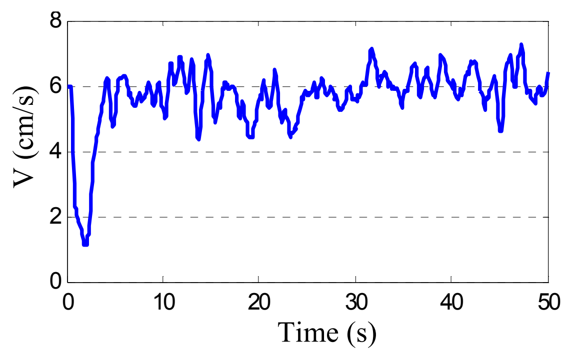

Figure 7. Velocity profile of a 300 cm drive with velocity control. Note that the velocity slows to zero while the follower waits for the leader to catch up, then increases to nominal and remains near nominal for the remainder.

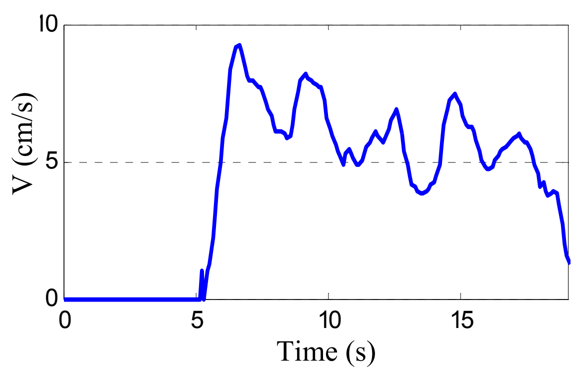

Figure 8. Velocity profile of a -80 cm crab with velocity control. As for drive, the velocity slows to zero while the follower waits for the leader to start moving away, then increases to and remains near nominal until stopping.

Figure 7 provides a velocity profile for a drive. The follower starts to drive before the leader, leading to a torque. This drops the velocity near zero until the leader begins to catch up. Then, the follower gradually increases speed until the torque is minimized and velocity reaches nominal. Once at nominal velocity, the velocity oscillates slightly compensating for small changes in force due to slip and variation in leader velocity. A crab velocity profile, with similar results, is shown in Figure 8.

No failures of formation occurred during cooperative transport experiments. By using velocity control and ensuring the control begins before forces and torques build up due to partner motion, the rigid formation following is very robust and maintains safe forces on the manipulator arms.

4.2 Cooperative Alignment and Beam Placement

The goal of maintaining formation is to allow the team to successfully and reliably perform the desired task: placement of the beam component into the structure. To demonstrate the reliability of our approach, the team repeated alignment and beam placement multiple times. Any failure was recorded, including failures due to high arm forces or torques (formation failure), failure to align accurately enough for proper beam placement (alignment failure), and failure to place a beam properly due to an unrecognized poor alignment by rover or arm (placement failure). The team is positioned such that at least one rover can see the structure and is placed in formation with minimal forces and torques at a variable relative angle to the structure. Reference forces and torques are set to the initial condition.

Table III shows number of alignment attempts listed and number of failures by type. The first series of experiments were run using a previous alignment procedure that did not autonomously correct in the case of both robots simultaneously losing sight of the structure, and coarse alignment (only one robot seeing the structure) used a fixed angle rather than local angle. In the second series, the described approach is used. A series of photographs illustrating an alignment and placement is shown in Figure 9. In this example, only SRR can initially see the structure, and the team proceeds through coarse alignment crab, drive, and turn. After one iteration, both rovers can see the structure and the team transitions to fine alignment.

Table III: Team Alignment and Beam Placement Results

Number of Alignments	Formation Failures	Alignment Failures	Placement Failures
12	0	1	0
5	0	0	0

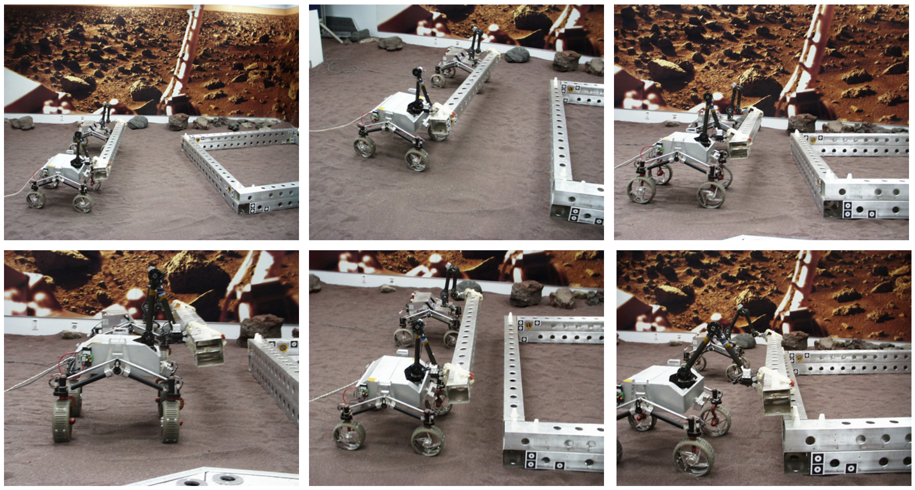

Figure 9. Alignment example. *Top left*: Initially only SRR sees the structure. *Top center*: After initial crab and drive based on SRR's observations. *Top right*: Fine alignment in progress after Ackermann turn based on both robots' observations.. *Bottom left*: Final crab to align for placement. *Bottom center*: Aligned for placement. *Bottom right*: Placing beam.

In the first set there was one failure due to both robots simultaneously losing sight of the structure. No autonomous correction was implemented and the operator corrected the error. There were no failures during the runs of the improved alignment procedure, including a case in which both robots lost sight of the structure and compensated. The iterative process allows correction of errors and multiple checks for validating alignment, making the cooperative transport and manipulation very robust.

5. Conclusions and Future Work

The Robotic Construction Crew performs cooperative transport and precision manipulation of long rigid beams in the context of a construction task. To place these beams precisely into a structure, the team must align accurately with the structure while remaining in formation. This process iteratively aligns range, offset, and heading. To maintain formation during transport, force-torque feedback is used to adjust velocity. RCC is robust to variable initial conditions, changes in the amount and quality of information available, synchronization errors and temporary communication loss, motion errors due to slippage, and minor driving or arm kinematics errors.

In current and future work, we will refine velocity control for formation following and investigate continual steering adjustments to compensate for formation errors not in the direction of motion. Additionally, force-torque feedback to maintain formations during turns will be included.

Acknowledgments

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. We thank Paul Schenker, Neville Marzwell and Eric Baumgartner for supporting the RCC effort, and Matthew Robinson for contributing to RCC development. We thank Brett Kennedy, Tony Ganino, Mike Garrett, and Lee Magnone for platform support.

References

Balch, T., and Arkin, R. C. (1998). Behavior-based formation control for multirobot teams. *IEEE Transactions on Robotics and Automation*, 14(6):926-939.

Borenstein, J. (2000). The Omnimate: A guidewire- and beacon-free AGV for highly reconfigurable applications. *International Journal of Production Research*, 38(9):1993-2010.

Brown, R. G. and Jennings, J. S. (1995) A pusher/steerer model for strongly cooperative mobile robot manipulation. *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, 3:562-568.

Carpin, S. and Parker, L. E. (2002). Cooperative leader following in a distributed multi-robot system. *Proceedings of IEEE Interernational Conference on Robotics and Automation*, 3:2994-3001.

Desay, J. P., Kumar, V., and Ostrowski, P. Control of Change in Formation for a Team of Mobile Robots. (1999). *Proceedings of IEEE International Conference on Robotics and Automation*, 2:1556-1561.

Huntsberger, T. L., Trebi-Ollennu, A., Aghazarian, H., Schenker, P., Pirjanian, P. and Nayar, H. D. (2004). Distributed Control of Multi-Robot Systems Engaged in Tightly Coupled Tasks. *Autonomous Robots*, 17:79-92.

Huntsberger, T., Pirjanian, P., Trebi-Ollennu, A., Nayar, H. D., Aghazarian, H., Ganino, A., Garrett, M., Joshi, S. S., and Schenker, P.S. (2003). CAMPOUT: A Control Architecture for Tightly Coupled Coordination of Multi-Robot Systems for Planetary Surface Exploration. *IEEE Transactions Systems, Man & Cybernetics, Part A*, 33(5): 550-559.

Mukaiyama, T., Kyunghwan, K., and Hori, Y. (1996). Implementation of cooperative manipulation using decentralized robust position/force control. *Proceedings of the 4th International Workshop on Advanced Motion Control*, 2:529-534.

NASA. (2004). Vision for Space Exploration.

Parker, L. E. (1994). ALLIANCE: an architecture for fault tolerant, cooperative control of heterogeneous mobile robots. *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2:776-683.

Qingguo, L. and Payandeh, S. (2003). Multi-agent cooperative manipulation with uncertainty: a neural net-based game theoretic approach. *Proceedings of IEEE International Conference on Robotics and Automation*, 3:3607-3612.

Rus, D., Donald, B., and Jennings, J. (1995). Moving furniture with teams of autonomous robots. *Proceedings of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 1:235-242.

Stroupe, A., Huntsberger, T., Okon, A., Aghazarian, H., and Robinson, M. (2005). Behavior-Based Multi-Robot Collaboration for Autonomous Construction Tasks. Submited to *IEEE International Conference on Robotics and Automation*.

Trebi-Ollennu, A., Das, H., Aghazarian, H., Ganino, A., Pirjanian, P., Huntsberger, T., and Schenker, P. (2002). Mars Rover Pair Cooperatively Transporting a Long Payload. *IEEE International Conference on Robotics and Automation*.